

SELF-MAINTAINING REAL-TIME DATA AGGREGATION

ABSTRACT OF THE DISCLOSURE

Multiple aggregation groups, which can be multiple partitions in an aggregated data table, are formed. Each group includes multiple aggregation records; each aggregation record includes an aggregation of values contained by a different subset of multiple database records. While an aggregation group is accessed by a single program thread during an aggregation group update transaction, no other threads are allowed to access that group. The aggregation groups are combined into a single table of aggregation records. Each of the multiple database records may correspond to an instance of an organizational activity and include a field having a value indicating the corresponding instance to be in one of several process states. Each aggregation group may further include time-sorted aggregation records, each time-sorted aggregation record containing an aggregation value for instances in one of the several process states during a time period associated with the time-sorted aggregation record. Aggregation records corresponding to instances completed outside of a preselected time window are deleted.